



PATENT

Case Docket No. MATH2.002AUS
Date: August 20, 2004

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s) : Troy Wray
App. No. : 10/817,765
Filed : April 2, 2004
For : ELECTROMAGNETIC FLOW
METER
Group Art Unit : 2855

I hereby certify that this correspondence and all marked attachments are being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on

August 20, 2004
(Date)

John M. Carson, Reg. No. 34,303

TRANSMITTAL LETTER

**Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450**

ATTENTION: MAIL STOP MISSING PARTS

Dear Sir:

In response to the Notice to File Missing Parts of Application Under 37 CFR 1.53(f), which was mailed by the Office on June 22, 2004, enclosed are:

- (X) An executed Declaration by Inventors.
- (X) A Power of Attorney Form and Copy of Assignment.
- (X) A certified copy of United Kingdom priority document numbers 0307637.9 and 0315612.2.
- (X) A Notice to File Missing Parts.
- (X) Return prepaid postcard.



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

CERTIFIED COPY OF PRIORITY DOCUMENT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., P.L.C. or PLC.

Registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 7 June 2004

BEST AVAILABLE COPY




1/77

03APR03 E797275-1 D02000
P01/7700 0.00-0307637.9

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference

IK/24868

-2 APR 2005

2. Patent application number

(The Patent Office will fill in this part)

0307637.9

3. Full name, address and postcode of the or of each applicant (underline all surnames)

ABB Limited
Stonefield Works
Oulton Road
Stone, Staffordshire
ST15 0RS

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

United Kingdom

8345282001 *IF*

4. Title of the invention

Electromagnetic Flow Meter

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

MATHYS & SQUIRE
100 Gray's Inn Road
London WC1X 8AL
United Kingdom

Patents ADP number (if you know it)

1081001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

YES

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description

21

Claim(s)

5

Abstract

Drawing(s)

8

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature

Mathys & Squire

Date

MATHYS & SQUIRE

2 April 2003

12. Name and daytime telephone number of person to contact in the United Kingdom

Ilya Kazi- 020 7830 0000

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

DUPLICATE

Electromagnetic Flow Meter

The present invention relates to electromagnetic flow meters. However, aspects of signal processing techniques disclosed herein may be more broadly applied.

- 5 The operating principles of Electromagnetic Flow Meters are well known, discussed for example in GB-A-2,380,798.

- 10 Where the sensing electrodes are in contact with the fluid, due to electrochemical or other effects, a DC potential is usually present across the electrodes even when there is no coil excitation, i.e. no field. That component of the signal is independent of the flow. This inhibits the ability to determine the flow in a static fashion. To overcome this some form of dynamic excitation to the coils is typically provided in order to generate a dynamic component at the electrodes that can be differentiated from the background DC (or slowly varying) bias signal. This
15 dynamic signal is normally pulsed DC or an AC signal.

- The bias signal will not generally be static. It may drift randomly with time, flow and temperature. In some applications, particularly for non-homogeneous fluids with inclusions, the signal may contain large amplitude decaying exponential
20 components, for example in paper pulp and slurry applications, as discrete charged particles occasionally touch the electrodes, changing the voltage instantaneously and this voltage then discharges exponentially to the (drifting) baseline.

- 25 The energy can be spread across a wide frequency range but most applications have significant low frequency noise and this is often more problematic due to the nature of the signal processing.

Therefore, one way to obtain flow measurements which are less susceptible to such effects might be to use a high frequency excitation, above most of the noise effects. It is found that a frequency above about 1 kHz would for most practical purposes exclude most bias effects. However, using such a frequency would
5 introduce its own problems. Most notably, the magnetic circuit of the flow meter is less stable at the higher frequencies. One factor that contributes to this is that the losses in the magnetic circuit, which become more significant at higher frequencies, are quite temperature dependent. Thus, for a given excitation, the field strength generated may vary unpredictably. This can be exacerbated
10 because the excitation coils are usually positioned outside a steel shell of the flow meter. A search coil can be used to measure the actual field generated but this adds significant cost and manufacturing difficulties.

The 'roll off' of the sensor head at these higher frequencies leads to an
15 uncertainty in the sensitivity of the sensor and to variations in the phase between the field (and hence the electrode signal) and the drive current. It is known to use a moderate frequency (e.g. 70 Hz) sinusoidal excitation and to demodulate the electrode signal synchronously with the excitation signal. Quite apart from other considerations, the phase angle at which the demodulation is (should be) carried
20 out is not constant and requires either manual or automatic adjustment.

To summarise the problems a lower drive frequency can give much better sensor stability but measurements are more easily corrupted by the bias drift and the effect of inclusions in the fluid. A further important problem is that a low excitation
25 frequency limits the rate at which new measurements can be updated - it gives a low flow measurement bandwidth. A higher frequency assists in distinguishing wanted signals from unwanted and also allows a more acceptable higher flow measurement bandwidth but the sensor characteristics will not be so precisely

defined and are less stable. The frequency chosen is therefore normally a compromise for a particular set of circumstances.

5 These problems have been previously addressed and it has been proposed in our earlier UK patent application no. 0116168.6 to excite a meter with two frequencies simultaneously or quasi-simultaneously. Each frequency component is separately extracted and a combined measurement obtained in such a way as to enhance the better properties of each measurement. A drawback is that measurement rate is still limited by the lowest frequency, as explained in more
10 detail below.

According to a first aspect the invention provides a method of obtaining a measure of flow from an electromagnetic flow meter comprising applying a composite excitation waveform comprising a plurality of frequency components;
15 receiving an output from the meter; and correlating sampled output data from the meter with a composite correlation waveform based on the composite excitation waveform to derive a measure of flow.

In this method, by taking the novel step of correlating with a composite waveform,
20 rather than the conventional method of correlating each frequency component with a sine wave, a number of advantages can be obtained. Although the composite waveform will effectively contain multiple frequency components, it has been found that it is not necessary to process each component separately and thus measurement rate is not limited by the lowest frequency component.
25 Furthermore, additional information is obtained from the combination of the components (the shape of the waveform) as well as the components themselves.

Preferably the meter is excited with a composite excitation waveform obtained by combining at least two frequency components; in this way the composite

waveform can easily be selected to contain frequency components which excite desired properties, for example a relatively high and relatively low frequency. The composite waveform will preferably not be completely sinusoidal. The composite waveform preferably is substantially devoid of sharp turning points (these are associated with numerous high frequencies which may behave unpredictably). The composite waveform is preferably digitally synthesised. The frequency components are preferably combined at pre-determined, preferably similar, for example within a range of minimum to maximum of 1 to 5 or less, preferably 1 to 3 or less, amplitudes. The amplitudes (and phases) may be chosen to enhance the prominence of turning points in the waveform.

The composite waveform preferably has a plurality of turning points within a predetermined measurement window period; it is found that multiple turning points or "features" may assist in obtaining accuracy.

More than two components may be used and provision of three or four components may enhance accuracy, by giving more features to the waveform. However, it is preferable that the number of components is limited, for example to less than about 10, preferably less than 5 components and/or that the highest frequency component with significant amplitude has a frequency no greater than about 1kHz. We have found that diminishing returns are achieved with greater than about 5 components, but in some cases a few more components may be used.

Whilst an arbitrary waveform can theoretically be approximated by an infinite number of sine and cosine waveforms, such approximations tend to require a large number of harmonics to converge accurately and this implies a large number of high frequencies which would behave unpredictably in a real meter. By limiting the number of components, most preferably by synthesising the

composite waveform from a discrete set of components, each having a chosen amplitude, the effects of the individual components and the overall response becomes more predictable.

- 5 Although the higher frequency may be an integral multiple of the lower frequency, it is not necessary for the frequencies to be harmonically related. Even when two frequencies are integral multiples of each other, it is not necessary to include intervening harmonics. Thus the composite waveform may comprise at least two frequencies which are not harmonically related to each other. Alternatively, the
- 10 composite waveform may comprise two frequencies of which a higher frequency is a harmonic of a lower frequency but in which at least one intervening harmonic of the lower frequency is substantially absent. Preferably, when the higher frequency is an odd harmonic, at least one intervening odd harmonic may be substantially absent and/or when the higher frequency is an even harmonic, at
- 15 least one intervening even harmonic may be substantially absent.

An advantageous feature is that correlation may be performed over a window which does not contain an integer number of periods of all (or any) frequency components. The window may be shorter than the period of the lowest

20 frequency component; this enables measurement to be obtained with a higher bandwidth than the lowest frequency.

Correlating preferably comprises determining a measure of closeness of fit. A factor may be derived from said correlating and a measure of flow may be

25 derived from the factor. A measure of a trend within a correlation window may be determined and a measure of measurement accuracy may be determined from the measure of trend.

With traditional methods of measurement, it is difficult to reduce sensitivity to trends in the baseline reading, e.g. due to zero offsets. Pursuant to the invention, it has been appreciated that in fact, particularly with the techniques disclosed herein (but other techniques which give multiple measurement may be used), trends (particularly baseline trends) in the meter data can usefully be measured and can be used to determine or improve meter accuracy.

According to a further aspect, the invention provides a method of processing the output of an electromagnetic flow meter comprising correlating sampled output data with a signal to determine a measure of flow and a measure of a trend within the data over a measurement period.

The trend may provide data about the meter or may simply account for interference during the measurement period.

A measure of a linear trend may be determined. Additionally or alternatively a measure of an exponential trend (e.g. a decaying exponential) may be determined. Advantageously an exponential or linear trend may be approximated by a polynomial. In a preferred arrangement a low order polynomial trend, preferably second or third order, is fitted to the data - this can effectively remove linear or exponential trends without interfering substantially with fitting of the data. The sampled output data may be fitted with reference data comprising an expected signal multiplied by an unknown factor and a trend of unknown gradient (or polynomial coefficients) to determine values for the factor and for the gradient (or polynomial) which gives the best fit according to a predetermined fitting method, preferably a least squares fitting method. Using higher order polynomials to model trends leads to diminishing returns. Preferably a measure of flow is determined from the value of the factor. Preferably a measure of confidence or accuracy is determined from the gradient.

The predetermined fitting method may comprise determining a sum of squares of differences between the reference data and sampled output data. The reference data may include a constant unknown offset. A value for the offset may be determined to improve the fit. However, advantageously, the offset value may be
5 determined but the measure of accuracy or confidence is based primarily on the value of the gradient. Calibration parameters for the meter may be stored, the method further comprising adjusting at least one calibration parameter for the meter based on a plurality of measurements of trend or accuracy or confidence.

10 According to a further aspect, the invention may provide a method of operating a flow meter comprising storing calibration parameters for the meter and adjusting at least one stored parameter as successive flow measurements are derived based on a plurality of measures of accuracy of the flow measurements.

15 Preferably at least one stored calibration parameter is adjusted to reduce a measure of trend or to improve a measure of accuracy or confidence based on successive measurements.

20 Calibration parameters may include at least one of phase and amplitude response. Preferably calibration parameters are stored for multiple frequency components. At least one calibration parameter for a first frequency component may be adjusted in response to measurements whereas at least one corresponding calibration parameter for a second frequency component may be retained at an initial (e.g. factory calibration) setting. The first frequency
25 component may have a higher frequency than the second frequency component. Preferably calibration parameters for the lowest frequency component are all kept constant.

Weighted filtering may be performed on measurements of flow based on the measure of trend or accuracy or confidence.

5 An embodiment of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:-

Figure 1 shows an example time window with two components for the case of two frequency components and a good fit;

Figure 2 shows another example of a good fit with four frequency components;

10 Figure 3 shows an example similar to Fig. 2 but exhibiting a poor fit;

Figure 4 illustrates removal of a trend from data;

Figure 5 shows the expected "pure" signal from the electrodes of an electromagnetic flow meter of an embodiment;

Figure 6 is a schematic diagram of apparatus according to an embodiment;

15 Figure 7 illustrates error processing using an SSE approach;

Figure 8 illustrates an enhanced method for determining errors;

Figure 9 shows a flow signal as determined by an LSE fit algorithm;

Figure 10 shows flow after filtering;

Figure 11 shows a flow profile;

20 Figure 12 shows the added noise zoomed in so the total window is about 0.2 seconds long;

Figure 13 shows derived flow values;

Figure 14 shows the error estimates for each of several windows;

Figure 15 shows noise added (over a 6 second window);

25

The method underlying the invention and a background method will first be explained.

As noted above, single frequency measurement in an electromagnetic flow meter is a compromise. To assist in understanding the invention, we will first explain in more detail an earlier technique which applies two frequencies simultaneously and then resolves the two components of the electrode signal separately. This
5 gives some of the benefits of each frequency. Such a system is disclosed in our earlier UK patent application no. 0116168.6.

As a specific example we consider a system using 10 Hz and 70 Hz excitation simultaneously. At 10 Hz, the sensor characteristics can be assumed to be stable
10 but the results will be influenced by noise. At 70 Hz, the signal is easier to recover but since the sensor characteristics are less well defined, the flow signal is correspondingly less well defined.

A basic system could use the zero measurement from the 10 Hz 'channel' as the
15 reference zero. The 70 Hz channel could then be used to obtain a measurement with the benefits of the higher frequency excitation and the zero would be 'anchored' to the 10 Hz value.

Doing the 10 Hz correlation in the presence of the 70 Hz component is not
20 problematic; the higher frequency is an integral multiple of the lower frequency (this is a constraint of the method) and an integer number of cycles of an interfering frequency in the correlation window does not cause a problem. However, performing the 70 Hz correlation is problematic. It is necessary to arrange that the correlation window to contain an integer number of both the 70
25 Hz and the 10 Hz signals in order to get an accurate result. This means that the window length for the 70Hz signal should be at least 0.1 seconds and hence measurement bandwidth is limited to that of the lower frequency, 10 Hz. Thus, although some noise may have been reduced and stability gained, the measurement bandwidth is still limited by the lower frequency.

Our earlier application in fact discloses a more sophisticated technique with complementary high and low pass filters but again the measurement bandwidth is limited by the lower frequency.

5

In conventional multi-frequency analysis, the fact that superposition applies (or is assumed to apply) is a useful basis for simplifying calculations; it allows the components to be resolved independently even if applied simultaneously as if the others weren't there (although this is subject to some limitations).

10

In the present invention, a composite waveform is applied, and this can be considered to be applying multiple excitation frequencies simultaneously (the coil excitation is the sum of a number of discrete sine waves). An important difference between embodiments of the invention and prior multi-frequency techniques is that the electrode signal is analysed for all components simultaneously by correlating the actual electrode signal with the expected electrode signal, rather than by detecting each frequency component separately.

15

To assist in explaining the invention an example window in time is shown in Figure 1.

20

The "curve 1" line shows the expected electrode signal, the "curve 2" noisy data is the actual electrode signal. The magnitude of the correlation is then the magnitude of the (curve 1) curve that gives the best fit to the (curve 2) data. In the figure, the fit is already optimal.

25

Note that in this example, there are two frequency components and both have an integer number of cycles in the window. In this case, resolving the two frequencies separately would give similar results to treating the composite

waveform as one signal. In this case, the inventive method may offer an advantage in the processing but it is noted that the measurement should be similar to that obtained by treating the components separately.

- 5 Figure 2 shows another example where the inventive method gives a significant advantage. In this example, there are four frequency components, two of them containing less than one complete cycle in the window.

10 Provided the expected signal is known then one of the great benefits of this technique is that it does not depend on any or all of the components having an integer number of cycles within the window.

15 Conceptually, we are looking for the 'pattern' (curve 1) in the data (curve 2). In principle, the more complex curve 1, the more distinct a pattern we are looking for.

20 Relative to the window above, the four frequency components have frequencies 1, 2, 0.5 and 0.25. The four frequencies are 40 Hz (1 cycle in the window) 80 Hz, 20 Hz and 10 Hz. In this example, the frequencies are harmonically related. In practice, using a digital signal processor (DSP), it may often be convenient to produce frequencies which are integer *sub*-multiples of a common frequency and this will often mean that the higher frequencies are harmonics of lower frequencies but it is not necessary for the components to be related.

- 25 The ability to fit without requiring a whole number of cycles of all of the frequencies provides significant advantages.

It is advantageous from a 'distinctiveness' point of view to have several turning points in the waveform within the window, hence the choice of frequencies and window length.

- 5 Ignoring zero offsets and noise for the moment then the expected signal from the electrodes should be of the form shown in Figure 5 and the amplitude should be directly proportional to the value of flow. For example, if Figure 5 represents the received signal at a flow of 1 m/s then a flow of 0.5 m/s should give half the amplitude, a flow of 0 m/s would give a flat line and a flow of -1 m/s would give the same signal as Figure 5 but inverted.
- 10

Note that this waveform shape is not intended to be fixed or indeed optimal. As noted above, the example happens to use harmonically related frequencies. This is not necessary and may not be the most advantageous arrangement. In any event it should be recognised that the pattern being matched will change from window to window.

15

In this case of no zero offsets, the waveform can be fitted using simple and well known Least Squares; the amplitude of the red curve is the one that minimises the sum of the square of the errors between the two sets of the data.

20

If our expected signal is P_i and our received signal is represented by Y_i then we are going to find the factor F such that:

$\sum (Y_i - F P_i)^2$ is a minimum. In this case, our only degree of freedom is F and we simply choose F to minimise the SSE.

However, working with real world data, we have recognised that the signal we are fitting is continually being corrupted by a varying offset and the offset may vary within the window we are fitting.

- 5 Two extra degrees of freedom have been introduced, namely a slope and an offset. Now the SSE equation that we want to minimise becomes:

$$\sum (Y_i - F \cdot P - M \cdot i - C)$$

- 10 Referring to Figure 4, curve 1 represents the incoming data with a clear trend on it. Curve 2 represents the same data after fitting using the above equation to remove the trend.

- 15 This idea will seem deceptively simple on first inspection. Indeed the technique could be applied to single frequency excitation systems to achieve a certain amount of robustness to the incoming signal having a superimposed trend.

- 20 However, trying to apply the technique to a multi-frequency system where the signals are analysed separately would be problematic. It would only work provided the expected signal does not have a trend within the window! This means that there could only be whole number of cycles of any applied frequencies within the window.

- 25 This technique works so well because we are fitting the composite signal directly and simultaneously. Where it demonstrates its power very effectively is on the 'tails' of exponentials. As mentioned earlier, the electrode signal could have a trend that is nothing to do with the flow. Most frequently this would be either a lowish frequency drift or the tail of an exponential. In the above graph, the magenta curve is the measured data. It has a negative going trend that would

lead to a poor fit with the red reference curve. However, by doing the fit with the slope and offset degrees of freedom, we obtain a good fit (the blue data has the trend removed).

- 5 It should be noted that this is very different to just removing any trend before the fit because even the expected signal has an inherent trend by virtue of not having a whole number of cycles of the waveforms in the window and so simply attempting to remove a trend would corrupt the data.
- 10 In the mathematics, the scale factor that gives the best fit (peak to trough) is the flow signal and the slope and offset can be thought of as the best trend that can be applied to the signal to make the peaks and troughs coincide.

The scale factor is F , the slope is M and the offset is C . The offset C is thrown away, it offers little informational value. However, M gives surprisingly useful information. Although at first sight M would be expected to be no longer needed once the fit is achieved, we have found that M in fact gives a very useful indication of the amount of corruption seen by the system during the measurement window.

- 20 The difference between a poor fit (above) and a good fit (the previous two graphs) is very apparent by eye. However, mathematically, the quality can be much harder to judge in a reliable but computationally simple manner. This is because there is a moderate level of random noise on the data anyway and even
- 25 a perfect fit has a moderate rms error. A further aspect of this invention provides a method to determine the quality of fit. This is done in a way that can be seen mathematically to enhance the distinction between Figures 2 and 3.

We now consider a practical system. It is a sampled data system. The sensor is being excited with 10, 20, 40 and 80 Hz (these may be conveniently generated e.g. as submultiples of a 20 kHz or 20.48 kHz DSP sampling clock frequency with 250 or 256 samples respectively for the 80 Hz signal and correspondingly more
5 for the lower frequencies). The acquisition system is acquiring samples continuously but the processing is done on a window of data at a time.

The window length is one cycle of 40 Hz (25 mS) but the window is moved along half a cycle at a time so the windows overlap (it is an advantageous but optional
10 feature that windows overlap, although it is not necessary to use exactly half a window increment). While this does not directly increase the bandwidth, it does give a higher data rate making subsequent filtering easier.

From the result of each window is a flow value F , a slope value M and an error estimate for the fit of the waveform within that window. Sometimes, the M and error estimates will suggest that the value of F is accurate and can be trusted highly. Other times M and/or the error estimate will suggest that the value of F could well be corrupt because there was excessive trend in the data or the fit just was not very good.
20

This confidence value can be used to 'weight' the data for filtering purposes. Again, this filtering has proven to be very effective on real world data. Different filtering and weighting algorithms may be used, the choice may depend on application.

25 A simple algorithm is to hold the last value if the confidence is too low. A monitor can be provided to ensure that the algorithm adapts if the overall reading confidence for all the windows is low.

In a more advanced algorithm, trend analysis is used to give a better guess of the value if the confidence in the current reading is low. In this respect, the filter could behave like a Kalman filter.

- 5 Equally, the M value / error estimate can be used to give a SEVAL like output e.g. clear, blurred, dazzled as proposed by M. P. Henry et al.

The ability to fit the incoming signal and remove a trend is a unique feature of this kind of system, which can be applied to other measurements, but would not have
10 an equivalent with a single frequency excitation.

The system depends on knowing the amplitude and phase of the incoming components. This information represents the frequency response of the sensor. However, we appreciate that this is not necessarily constant. An algorithm will
15 now be described that can automatically adapt to changing sensor characteristics.

An outer loop is responsible for updating the matrix of coefficients representing the frequency response. By suitable averaging of the residual errors of each
20 window, any error in one of the coefficients becomes apparent.

The fit that gives the best possible fit while keeping the amplitude ratio and phase of the components fixed is used. However, if the fit could be improved by allowing the matrix coefficients to be adjusted (preferably by a small amount below a
25 threshold) and the degree of adjustment is consistent over a number of windows then the adjustment is made.

Preferably the lower frequency, here 10 Hz coefficients are kept fixed (but in some cases may be subject to a very low limit of adjustment). Limits can be

placed on the amount by which the other components can be adjusted. Limits may include a limit on the maximum total deviation of a given coefficient from a starting value and/or a maximum adjustment of a given coefficient in a given adjustment interval and/or a maximum total or weighted combination of adjustment of coefficients and/or a combination of the preceding limits.

A method of determining measurement error will now be explained, referring back to Figure 1. It is clear by eye that there would be no better fit than the one shown. However, the SSE (Sum of the Squared Errors) is still significant because of the (relatively) high frequency noise on the data being fitted. Now referring to Figure 3, by eye, it is clear that the fit is not very good. However, the SSE is not very different from that for Figure 1. For illustration, the SSE for figure 1 is 0.121 while for Figure 3 the SSE is 0.269, a little over a factor of 2. Yet the fit is very much worse. This illustrates the normal difficulty in measuring accuracy of measurement.

Since we want to use the measurement error as an aid to weighted filtering, we want there to be more 'contrast' (by which we mean a mathematically identifiable distinction) between the error values for Figure 1 and Figure 3. We have found that we can achieve this by substantially ignoring the high frequency residuals and considering primarily the low frequency residuals. Conceptually therefore we could filter the residuals with a low pass filter before performing the SSE calculations. Alternatively, we could do an FFT of the residuals and measure only components below a certain frequency. Both of these methods may be effective, but performing an FFT at least is computationally intensive.

A computationally simple but highly effective method for obtaining much more contrast will be described. This is based on the recognition that the residual errors in figure 3 that we want to measure are low frequency and that, as a result, the error spends a long time on the same side of the curve (ie either positive or

negative). In contrast, the high frequency residuals average to zero over a very short interval.

The novel method is to divide the window into a plurality of intervals, preferably at least about 4, preferably no more than about 20, typically 5-10, here 8 intervals. For each of these intervals, the residual is summed algebraically. Referring to figure 1 it is clear that this algebraic sum would be close to zero over any of the intervals whereas for figure 3 it would be far from zero (either large and negative or large and positive) over most of the intervals.

To get a final error figure for the window, the absolute value of each of the 8 interval errors is summed. This is because the summation within any interval could be positive or negative and an algebraic sum would largely cancel out.

There is thus provided a method of determining error between a signal and a corresponding reference signal, each signal comprising a first plurality of sampled values in a sampling interval, the method comprising:

defining a second plurality of sub intervals, each sub interval containing a respective subset of the first plurality of sampled values;

determining an algebraic sum for each sub interval of the difference between signal sample values and reference signal sample values;

determining an absolute sum of the algebraic sums for each sub interval as an error estimate for the sampling interval.

The calculation may include other filtering or weighting and the error estimate may be further processed, but the important point is that the step of algebraic summing comprises a calculation or operation in which differences of opposite sign substantially cancel and the step of absolute summing comprises an operation or calculation in which the algebraic sums add constructively. In place

of an absolute sum, a sum of square might be used - this would increase the effect of a single large error on the end result.

Although the units are different, the error using this technique is 2.203 and
5 48.033 for figures 1 and 3 respectively; a much higher contrast (20:1 instead of 2:1)

Dividing the window into 8 intervals is very appropriate in this example because
of the frequencies present (ie ratios of 1, 2, 4 and 8). If different frequencies were
10 chosen then these intervals could be optimised accordingly.

Figure 7 shows the SSE approach. The graph represents 6 seconds worth of 25
ms windows. Although there is high contrast between the peaks and the
'baseline', these peaks only represent the very worst readings.
15

Figure 8 shows the enhanced method for determining errors. Here, the baseline
is much lower and much better contrast exists between good and bad readings.

On these two graphs, every point represents a processed window. Figure 9
20 shows the flow signal as determined by the LSE fit algorithm. The final stage of
the processing now is to use the error measurement from graph 8 to filter the
data in graph 9. Since it is clear that the high peaks in figure 8 represent points
with high measurement errors, we can 'filter out' these points with a suitable
algorithm.

25

Figure 10 shows the flow after filtering.

Figure 11 shows the flow profile used for the examples in this example. Figure 15 shows the noise added. Although the results are partially simulated, the noise added to the data is from a real flow meter in a real application.

- 5 Figure 12 shows the added noise zoomed in (the total window is about 0.2 seconds long).

Figure 13 shows the derived flow values for approximately the same interval. The intervals in Fig 12 show the window numbers and correspond to the x axis values
10 of figure 13.

Figure 14 shows the error estimates for each of the windows, demonstrating that we have a sensitive indicator of when the flow estimate is not perfect. Note that the flow reading is sometimes good even though the error indicator is high. This is logical because the error indicator is really indicating the uncertainty in the flow
15 reading, at least in a qualitative way.

Apparatus embodying the method will now be described with reference to Figure 6.

- 20 The Microprocessor 10 maintains a 'time' variable for the purpose of calculating drive signals and expected signals. In this embodiment, the microprocessor calculates a signal that is the sum of four predetermined amplitude (here equal amplitude), distinct frequency sinewaves. This signal is applied to the drive coil in the sensor 3 via power amplifier 2. The electrode signals are amplified 4 and
25 converted into discrete samples by ADC 5. Software modules running on the microprocessor implements the method described.

The mathematical drive signal is transformed by the frequency response coefficients in transformation module 9 in order to determine the expected signal

back from the sensor. In other words, knowing the signal applied to the drive coil and knowing the transfer function of the sensor, this module calculates the waveform expected back from the sensor in the absence of perturbations.

- 5 Correlation module 6 performs the correlation or LSE fitting. The results of this fitting are at least a flow signal and a 'confidence' value that can be used to assist in the weighted filter module 7. As an enhancement, the results from the correlation module 6 can also be used by an adaptive adjustment module 8. Such a module, if present, monitors the quality of fit over a period of time and
- 10 decides whether the coefficients used by transformation module 9 need adjustment, subject to predetermined constraints.

- Embodiments of the invention provide one or more of the following novel features: Using multiple frequencies simultaneously, fitting multiple frequencies
- 15 simultaneously, weighted filtering and adjusting coefficients dynamically.

Modifications of detail may be made and features disclosed herein may be provided independently or in other combinations.

Claims

1. A method of obtaining a measure of flow from an electromagnetic flow meter having at least one field generating coil and potential sensing electrodes, the method comprising:

5 applying to said at least one field generating coil a composite excitation waveform comprising at least two frequency components and wherein a plurality of frequency components are simultaneously present for at least a portion of the waveform;

10 sampling a potential sensed by the potential sensing electrodes to produce a sampled output;

 correlating the sampled output with a composite waveform based on the composite excitation waveform;

 deriving a measure of flow based on the result of said correlating.

15 2. A method according to Claim 1 wherein the waveform comprises a set of components, each having a pre-specified amplitude and frequency.

3. A method according to any preceding claim wherein the composite waveform is non-sinusoidal.

20 4. A method according to any preceding claim wherein the composite waveform is digitally synthesised.

25 5. A method according to any claim as dependent on Claim 1 wherein correlation is performed over a window which does not contain an integer number of periods of all frequency components.

6. A method according to Claim 5 wherein the window is shorter than the period of the lowest frequency component.

7. A method according to any claim as dependent on Claim 1 wherein correlating comprises determining a measure of closeness of fit.
- 5 8. A method according to any claim as dependent on Claim 1 wherein a factor is derived from said correlating and a measure of flow is derived from the factor.
9. A method according to any claim as dependent on Claim 1 wherein a
10 measure of a trend within a correlation window is determined.
10. A method according to Claim 9 wherein the measure of trend is subtracted from the output sample.
- 15 11. A method according to Claim 9 or 10 wherein a measure of measurement accuracy is determined from the measure of trend.
12. A method of processing a signal, preferably the output of an electromagnetic flow meter, comprising correlating sampled output data of the
20 signal with a reference signal to determine both a measure of a physical quantity, preferably flow, and a measure of a trend within the data over a measurement period.
13. A method according to Claim 10 or 11 wherein a measure of a linear trend
25 is determined.
14. A method according to Claim 12 wherein the sampled output data is fitted with reference data comprising an expected signal multiplied by an unknown factor and a trend of unknown gradient to determine values for the factor and for

the gradient which gives the best fit according to a predetermined fitting method, wherein preferably a measure of physical quantity (preferably flow) is determined from the value of the factor and wherein preferably a measure of confidence or accuracy is determined from the gradient.

5

15. A method according to Claim 14 wherein the predetermined fitting method comprises determining a sum of squares of differences between the reference data and sampled output data.

10 16. A method according to Claim 14 or 15 wherein the reference data includes a constant unknown offset.

17. A method according to Claim 16 wherein a value for the offset is determined to improve the fit.

15

18. A method according to Claim 17 wherein the offset value is determined but the measure of accuracy or confidence is based primarily on the value of the gradient.

20 19. A method according to any of Claims 9 to 18 wherein calibration parameters for the meter are stored, the method further comprising adjusting at least one calibration parameter for the meter based on a plurality of measurements of trend or accuracy or confidence.

25 20. A method of operating a flow meter comprising storing calibration parameters for the meter and adjusting at least one stored parameter as successive flow measurements are derived based on a plurality of measures of accuracy of the flow measurements.

21. A method according to Claim 20 wherein at least one stored calibration parameter is adjusted to reduce a measure of trend or to improve a measure of accuracy or confidence based on successive measurements.

5 22. A method according to Claim 19, 20 or 21 wherein calibration parameters include at least one of phase and amplitude response.

23. A method according to Claim 22 wherein calibration parameters are stored for multiple frequency components.

10

24. A method according to Claim 23 wherein at least one calibration parameter for a first frequency component is adjusted in response to measurements and at least one corresponding calibration parameter for a second frequency component is retained at an initial setting.

15

25. A method according to Claim 24 wherein the first frequency component has a higher frequency than the second frequency component.

26. A method according to Claim 25 wherein calibration parameters for the
20 lowest frequency component are all kept constant.

27. A method according to any of Claims 9 to 26 wherein weighted filtering is performed on measurements of flow based on the measure of trend or accuracy or confidence.

25

28. A method of determining error between a signal and a corresponding reference signal, each signal comprising a first plurality of sampled values in a sampling interval, the method comprising:

defining a second plurality of sub intervals, each sub interval containing a respective subset of the first plurality of sampled values;

determining an algebraic sum for each sub interval of the difference between signal sample values and reference signal sample values;

- 5 determining an absolute sum of the algebraic sums for each sub interval as an error estimate for the sampling interval.

29. A method substantially as herein described.

- 10 30. Apparatus substantially as herein described or as illustrated in Fig. 6 of the accompanying drawings.




Figure 1




Figure 2






Figure 3




Figure 4






Figure 5




Figure 11



.

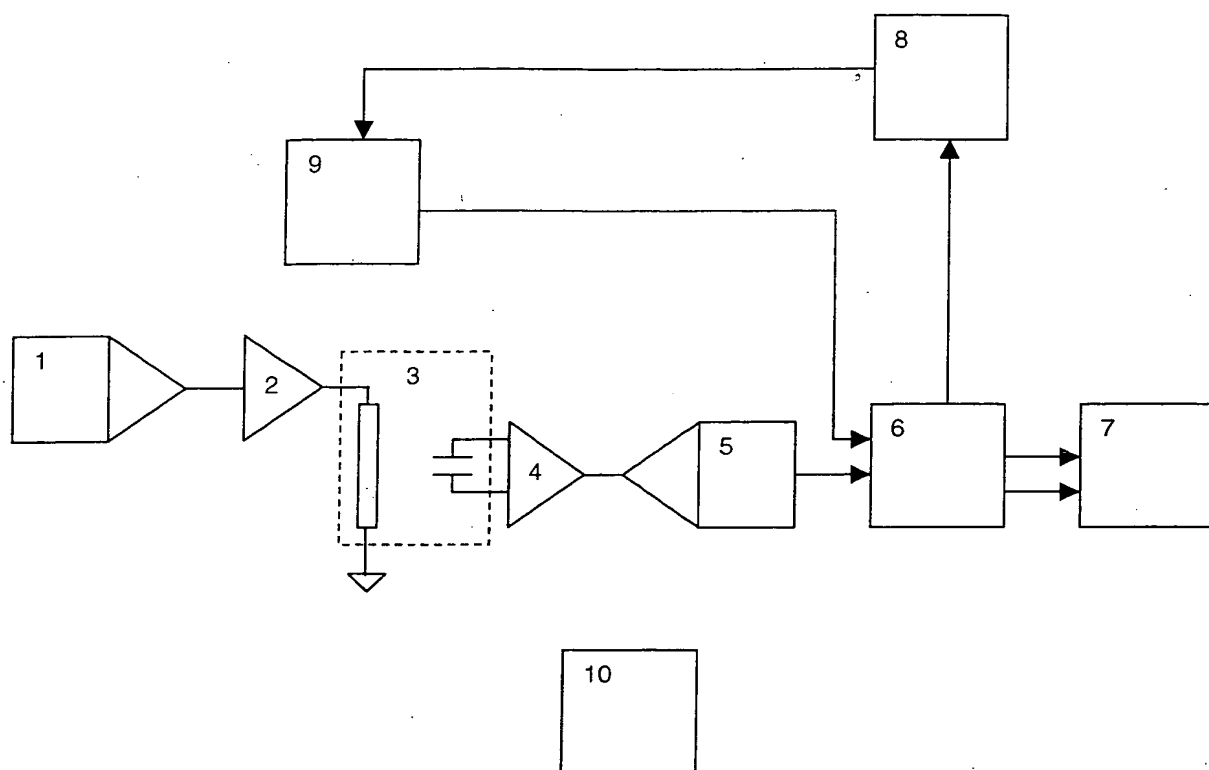


Figure 6






Figure 7




Figure 8






Figure 9




Figure 10






Figure 12




Figure 13




Figure 14






Figure 15

